arXiv:1701.03077v3 [cs.CV] 9 Mar 2017

A More General Robust Loss Function

Jonathan T. Barron
barron@google.com

Abstract

We present a two-parameter loss function which can be viewed as a generalization of many popular loss functions used in robust statistics: the Cauchy/Lorentzian, Geman-McClure, Welsch, and generalized Charbonnier loss functions (and by transitivity the L2, L1, L1-L2, and pseudo-Huber/Charbonnier loss functions). We describe and visualize this loss, and document several of its useful properties.

Many problems in statistics [8] and optimization [6] require *robustness* — that a model be insensitive to outliers. This idea is often used in parameter estimation tasks, where a non-robust loss function such as the L2 norm is replaced with some most robust alternative in the face of non-Gaussian noise. Practitioners, especially in the image processing and computer vision literature, have developed a large collection of different robust loss functions with different parametrizations and properties (some of which are summarized well in [2, 13]). These loss functions are often used within gradient-descent or second-order methods, or as part of M-estimation or some more specialized optimization approach. Unless the optimization strategy is co-designed with the loss being minimized, these losses are often "plug and play": only a loss and its gradient is necessary to integrate a new loss function into an existing system. When designing new models or experimenting with different design choices, practitioners often swap in different loss functions to see how they behave. In this paper we present a single loss function that is a superset of many of these common loss functions. A single continuous-valued parameter in our loss function can be set such our loss is exactly equal to several traditional loss functions, but can also be tuned arbitrarily to model a wider family of loss functions. As as result, this loss may be useful to practitioners wishing to easily and continuously explore a wide variety of robust loss functions.

1. Derivation

We will derive our loss function from the "generalized Charbonnier" loss function [12], which has recently become popular in some flow and depth estimation tasks that require robustness [4, 10]. The generalized Charbonnier loss builds upon the Charbonnier loss function [3], which is generally defined as:

$$f\left(x, c\right) = \sqrt{x^2 + c^2} \tag{1}$$

This loss sometimes is written in a reparameterized form:

$$f\left(x, c\right) = c\sqrt{\left(x/c\right)^2 + 1} \tag{2}$$

This form of the loss is sometimes referred to as "L1-L2" loss (as it behaves like quadratic loss near the origin and like absolute loss far from the origin) or a Pseudo-Huber loss (due to its resemblance to the classic Huber loss function [7]). The generalized Charbonnier loss function takes the Charbonnier loss and, instead of applying a square-root, raises the loss to an arbitrary power parameter α:

$$g\left(x, \alpha, c\right) = \left(\left(x/c\right)^2 + 1\right)^{\alpha/2} \tag{3}$$

Here we use a slightly different parametrization from [12] and use $\alpha/2$ as the exponent instead of just α. This makes the generalized Charbonnier somewhat easier to reason about with respect to standard loss functions: $g\left(x, 2, c\right)$ resembles L2 loss, $g\left(x, 1, c\right)$ resembles L1 loss, etc. We also omit the c scale factor in Equation 2, which gives us scale-invariance with respect to c:

$$g(2x, \alpha, 2c) = g(x, \alpha, c) \tag{4}$$

This allows us to view the c "padding" variable as a "scale" parameter, similar to other common robust loss functions.

But this formulation of the generalized Charbonnier still has several unintuitive properties: the loss is non-zero when $x = 0$ (assuming a non-zero values of c), and the curvature of the quadratic "bowl" near $x = 0$ varies as a function of c and α. We therefore construct a shifted and scaled version of Equation 3 that does not have these properties:

$$\frac{g\left(x, \alpha, c\right) - g\left(0, \alpha, c\right)}{c^2 g''\left(0, \alpha, c\right)} = \frac{1}{\alpha}\left(\left(\left(x/c\right)^2 + 1\right)^{\alpha/2} - 1\right) \tag{5}$$

This loss has the unfortunate side-effect of flattening out to 0 for large negative values of α, which we address

by modifying the $1/\alpha$ scale factor while preserving scale-invariance.

$$h(x, \alpha, c) = \frac{z(\alpha)}{\alpha}\left(\left(\frac{\left(x/c\right)^2}{z(\alpha)} + 1\right)^{\left(\alpha/2\right)} - 1\right) \tag{6}$$

$$z(\alpha) = \max\left(1, 2 - \alpha\right) \tag{7}$$

This loss resembles a normalized, centered, and scale-invariant version of the loss shown in Equation 2.

As is already well-known, L2 loss is a special case of the generalized Charbonnier, and this remains true in our reparametrization:

$$h\left(x, 2, c\right) = \frac{1}{2}\left(x/c\right)^2 \tag{8}$$

Though L1 loss is a special case of the generalized Charbonnier loss in it's traditional form, true L1 loss is not actually expressible in our normalized form in Equation 6 due to the division by c. But if we assume that x is much greater than c, we see that our normalized form approaches L1 loss:

$$h\left(x, 1, c\right) \approx \frac{|x|}{c} - 1 \quad \text{if } x \gg c \tag{9}$$

Unlike the generalized Charbonnier loss, our normalized loss can be shown to be a generalization of two other common loss functions. Though $h\left(x, 0, c\right)$ is undefined due to a division by zero, we can take the limit of $h\left(x, \alpha, c\right)$ as α approaches zero:

$$\lim_{\alpha \to 0} h\left(x, \alpha, c\right) = \log\left(\frac{1}{2}\left(x/c\right)^2 + 1\right) \tag{10}$$

Perhaps surprisingly, this yields the Cauchy (aka Lorentzian) loss function [1]. Cauchy loss is therefore a special case of our normalized loss, or equivalently, our normalized loss is not just a generalization of the Charbonnier, but is also a generalization of Cauchy loss.

Though this does not appear to be common practice, the power parameter α in a generalized Charbonnier can be set to negative values, and this is also true for our normalized variant. By setting $\alpha = -2$, our loss is equivalent to Geman-McClure loss [5]:

$$h\left(x, -2, c\right) = \frac{2\left(x/c\right)^2}{\left(x/c\right)^2 + 4} \tag{11}$$

And in the limit as α approaches negative infinity, our loss becomes Welsch [9] (aka Leclerc [11]) loss:

$$\lim_{\alpha \to -\infty} h\left(x, \alpha, c\right) = 1 - \exp\left(-\frac{1}{2}\left(x/c\right)^2\right) \tag{12}$$

The Welsch and Geman-McClure losses are therefore special cases of our loss, or equivalently, our loss function can be viewed as a generalization of the Welsch and Geman-McClure loss functions.

2. Loss Function

With this analysis in place, we can present our final loss function. Our loss is simply our normalized variant of the generalized Charbonnier loss, where we have introduced special cases to cover the otherwise-undefined limits as α approaches 0 and $-\infty$, as without these special cases the loss is technically undefined for these values.

$$\rho\left(x, \alpha, c\right) = \begin{cases} \log\left(\frac{1}{2}\left(x/c\right)^2 + 1\right) & \text{if } \alpha = 0 \\ 1 - \exp\left(-\frac{1}{2}\left(x/c\right)^2\right) & \text{if } \alpha = -\infty \\ \frac{z(\alpha)}{\alpha}\left(\left(\frac{\left(x/c\right)^2}{z(\alpha)} + 1\right)^{\left(\alpha/2\right)} - 1\right) & \text{otherwise} \end{cases} \tag{13}$$

$$z(\alpha) = \max\left(1, 2 - \alpha\right) \tag{14}$$

As we have shown, this loss function is a superset of the generalized Charbonnier loss function (and therefore the the Charbonnier / L1-L2 / pseudo-Huber, quadratic and absolute loss functions by transitivity) and is also a superset of the Cauchy / Lorentzian, Geman-McClure, and Welsch / Leclerc loss functions.

As a reference, we provide the derivative of $\rho\left(x, \alpha, c\right)$ with respect to x, for use in gradient-based optimization:

$$\frac{d\rho}{dx}\left(x, \alpha, c\right) = \begin{cases} \frac{2x}{x^2 + 2c^2} & \text{if } \alpha = 0 \\ \frac{x}{c^2}\exp\left(-\frac{1}{2}\left(x/c\right)^2\right) & \text{if } \alpha = -\infty \\ \frac{x}{c^2}\left(\frac{\left(x/c\right)^2}{z(a)} + 1\right)^{\left(\alpha/2 - 1\right)} & \text{otherwise} \end{cases} \tag{15}$$

This is also known as the influence function $\psi\left(x, \alpha, c\right)$ viewed through the framework of M-estimation.

As a reference for M-estimation, we also provide the weight function $w\left(x, \alpha, c\right)$ to be used during IRLS that corresponds to our loss:

$$\frac{1}{x}\frac{d\rho}{dx}\left(x, \alpha, c\right) = \begin{cases} \frac{2}{x^2 + 2c^2} & \text{if } \alpha = 0 \\ \frac{1}{c^2}\exp\left(-\frac{1}{2}\left(x/c\right)^2\right) & \text{if } \alpha = -\infty \\ \frac{1}{c^2}\left(\frac{\left(x/c\right)^2}{z(a)} + 1\right)^{\left(\alpha/2 - 1\right)} & \text{otherwise} \end{cases} \tag{16}$$

Let us now enumerate some properties of our loss function. Our loss is valid for all real values of α and for all real, positive values of c. The loss increases monotonically with the magnitude of x. The loss is scale invariant with respect to c and x:

$$\rho(2x, \alpha, 2c) = \rho(x, \alpha, c) \tag{17}$$

At the origin the loss is zero and the IRLS weight (when $\alpha \leq 2$) is 1.

$$\forall_{\alpha, c} \quad \rho\left(0, \alpha, c\right) = 0 \tag{18}$$

$$\forall_{\alpha,c} \quad \frac{1}{x}\frac{d\rho}{dx}\left(0,\alpha,c\right) = 1 \tag{19}$$

The roots of the second derivative of $\rho\left(x,\alpha,c\right)$ are:

$$x = \pm c\sqrt{\frac{a-2}{a-1}} \tag{20}$$

This tells us at what value of x the loss begins to re-descend. This point has a magnitude of c when $\alpha = -\infty$, and that magnitude increases as α increases. The root is undefined when $\alpha \geq 1$, from which we can infer that the loss is re-desending iff $\alpha < 1$.

For all values of x, α, and $c > 0$ the loss increases monotonically with α:

$$\frac{d\rho}{d\alpha}\left(x,\alpha,c\right) \geq 0 \tag{21}$$

This means that this loss can be annealed with respect to α in a graduated non-convexity framework. This property is due to our choice of $z(\alpha)$, as many other arbitrary clamping functions of α do not guarantee monotonicity.

For all values of α, when $|x|$ is small with respect to c the loss is well-approximated by a quadratic:

$$\rho\left(x,\alpha,c\right) \approx \frac{1}{2}\left({}^{x}/_{c}\right)^2 \quad \text{if } |x| < c \tag{22}$$

This approximation holds for ρ and its first and second derivatives. Because the second derivative of the loss is maximized at $x = 0$, this quadratic approximation tells us that the second derivative is bounded from above:

$$\frac{d^2\rho}{dx^2}(x,\alpha,c) \leq \frac{1}{c^2} \tag{23}$$

This property is useful when deriving approximate Jacobi preconditioners for optimization problems that minimize this loss.

When α is negative the loss approaches a constant as $|x|$ approaches infinity, which lets us provide an upper bound on the loss:

$$\forall_{x,c} \quad \rho\left(x,\alpha,c\right) \leq \frac{\alpha-2}{\alpha} \quad \text{if } \alpha < 0 \tag{24}$$

Additionally, when α is less than or equal to 1 we can provide an upper bound on the gradient of the loss:

$$\forall_{x,c} \quad \frac{d\rho}{dx}\left(x,\alpha,c\right) \leq \frac{1}{c}\left(\frac{\alpha-2}{\alpha-1}\right)^{\left(\frac{\alpha-1}{2}\right)} \quad \text{if } \alpha \leq 1 \tag{25}$$

A visualization of our loss and its derivative/influence and weight functions for different values of α can be seen in Figures 1, 2, 3, and 4. Because c only controls the scale of the loss on the x-axis, we do not vary c in our visualizations and instead annotate the x-axis of our plots in units of c.

3. Conclusion

We have presented a two-parameter loss function that generalizes many existing one-parameter robust loss functions: the Cauchy/Lorentzian, Geman-McClure, Welsch, and generalized Charbonnier, pseudo-Huber, L2, and L1 loss functions. We have presented the loss, gradient, and M-estimator weight as a reference, in addition to enumerating several convenient properties of the loss. By reducing a large discrete family of single-parameter loss functions into a single loss function with two continuous parameters, our loss function enables the convenient and continuous exploration of different robust loss functions, and and intuitive way to compare the relative effects of the members of these loss functions.

References

[1] M. J. Black and P. Anandan. The robust estimation of multiple motions: Parametric and piecewise-smooth flow fields. *CVIU*, 1996. 2

[2] M. J. Black and A. Rangarajan. The outlier process: Unifying line processes and robust statistics. *CVPR*, 1994. 1

[3] P. Charbonnier, L. Blanc-Feraud, G. Aubert, and M. Barlaud. Two deterministic half-quadratic regularization algorithms for computed imaging. *ICIP*, 1994. 1

[4] Q. Chen and V. Koltun. Fast mrf optimization with application to depth reconstruction. *CVPR*, 2014. 1

[5] S. Geman and D. E. McClure. Statistical methods for tomographic image reconstruction. *Bulletin de l'Institut international de statistique*, 1987. 2

[6] T. Hastie, R. Tibshirani, and M. Wainwright. *Statistical Learning with Sparsity: The Lasso and Generalizations*. Chapman and Hall/CRC, 2015. 1

[7] P. J. Huber. Robust estimation of a location parameter. *Annals of Mathematical Statistics*, 1964. 1

[8] P. J. Huber. *Robust Statistics*. Wiley, 1981. 1

[9] J. E. D. Jr and R. E. Welsch. Techniques for nonlinear least squares and robust regression. *Communications in Statistics-simulation and Computation*, 1978. 2

[10] P. Krähenbühl and V. Koltun. Efficient nonlocal regularization for optical flow. *ECCV*, 2012. 1

[11] Y. G. Leclerc. Constructing simple stable descriptions for image partitioning. *IJCV*, 1989. 2

[12] D. Sun, S. Roth, and M. J. Black. Secrets of optical flow estimation and their principles. *CVPR*, 2010. 1

[13] Z. Zhang. Parameter estimation techniques: A tutorial with application to conic fitting, 1995. 1

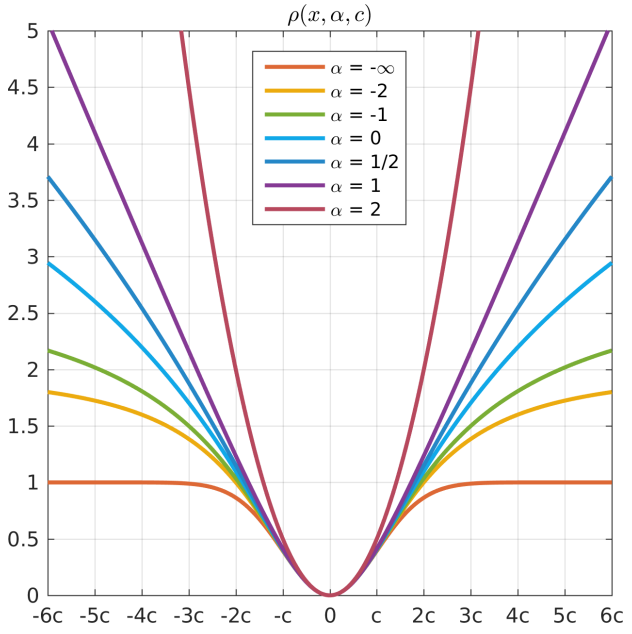

Figure 1: Our loss for different values of the power parameter α. When $\alpha = 2$ we have L2 loss, when $\alpha = 1$ we closely approximation L1 loss, when $\alpha = 0$ we have Cauchy/Lorentzian loss, when $\alpha = -2$ we have Geman-McClure loss, and as α approaches negative infinity we have Welsch loss.

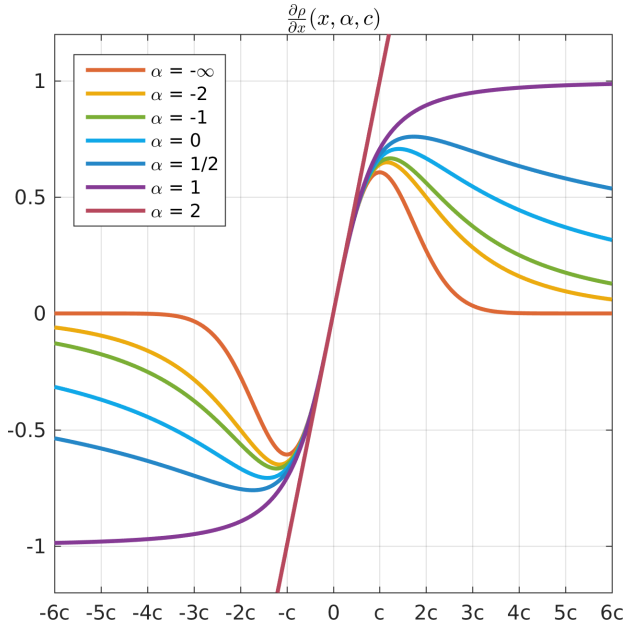

Figure 2: Our loss's gradient for different values of the power parameter α.

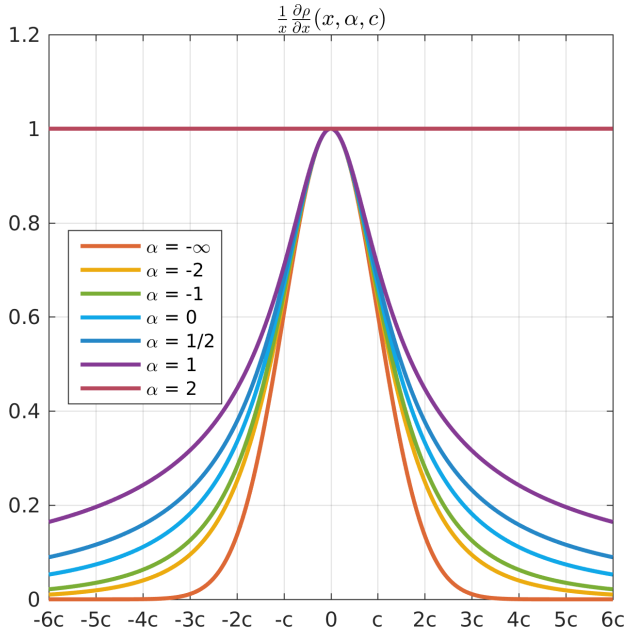

Figure 3: Our loss's IRLS weight for different values of the power parameter α.

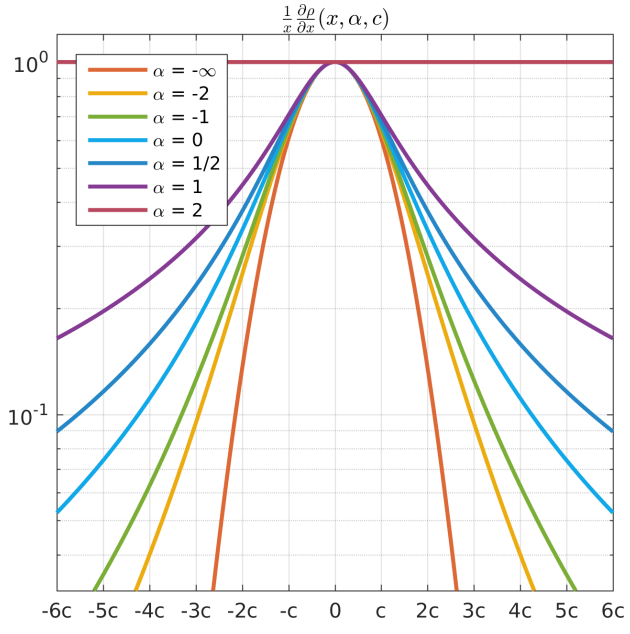

Figure 4: Our loss's IRLS weight for different values of the power parameter α, rendered using a logarithmic y-axis.